

EXHIBIT 5

UNITED STATES DISTRICT COURT
DISTRICT OF MASSACHUSETTS

SKYLINE SOFTWARE SYSTEMS, INC.,)
Plaintiff,) CIVIL ACTION NO.
v.) 06-10980-DPW
KEYHOLE, INC. and GOOGLE, INC.,)
Defendants.)

MEMORANDUM AND ORDER
November 16, 2006

The Plaintiff, Skyline Software Systems, Inc. ("Skyline"), alleges infringement of United States Patent No. 6,496,189 (the "'189 patent") by the Defendants, Keyhole, Inc. and Google, Inc ("Google"). The '189 patent describes a method and apparatus for streaming three-dimensional terrain data over a network, using data blocks arranged in a hierarchical data structure to improve performance. I have previously construed a set of terms related to this dispute, in a case docketed as Civil Action NO. 04-11129-DPW. *Skyline Software Systems v. Keyhole, Inc.*, 421 F.Supp.2d 371 (D. Mass. 2006).

Skyline has now broadened the dispute in this action to include infringement of patent claims other than 1 and 12, which it originally selected as "representative" of the dispute. Some of the additional claims contain disputed language not present in

Claim 1 or Claim 12, requiring me to engage in a second claim construction exercise. In addition, Defendants now dispute the meaning of terms found in Claims 1 and 12 and not construed in the previous *Markman* determination, arguing that these terms are pertinent to the interpretation of the additional claims.

Skyline makes a *quasi-res judicata* argument that the terms present in Claims 1 and 12 should not now be construed, but I find this argument unconvincing; unconstrued language in the original representative claims may now have more pertinence in the expanded context.

In the interests of judicial economy, I have directed the two cases be consolidated in this action and ordered the Plaintiff to file a comprehensive amended complaint for the case. In this Memorandum, I will construe terms the parties dispute, including additional terms found in Claims 1 or 12 where these terms may have relevance to the expanded claims.¹ I have been given no reason to revisit my earlier constructions. Consequently, all terms construed in the initial *Markman* Memorandum will continue to be construed throughout the claims of the patent as indicated there.

I. Claim Construction

The initial *Markman* memorandum, *Skyline Software Systems*,

¹ Attached as an Appendix to this Memorandum is a chart detailing the terms in dispute and the construction given to each.

421 F.Supp.2d at 374-77, fully laid out the legal framework for claim construction after *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005). I continue to follow that approach and incorporate it by reference here.

II. Disputed Terms

A. Disputed Terms in Claim 1 or Claim 12²

Claim 1 reads:

A method of **providing** data blocks describing three-dimensional terrain **to a renderer**, the data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels, the method comprising:

receiving from the renderer one or more coordinates in the terrain along with indication of a respective resolution level;

providing the renderer with a first data block which includes data corresponding to the one or more coordinates, from a local memory;

downloading from a remote server one or more additional data blocks at a resolution higher than the resolution level of the first block which includes data corresponding to the one or more coordinates **if the provided block from the local memory is not at the indicated resolution level.**

(col. 16, ll. 28-43) (emphasis supplied).

Claim 12 reads:

Apparatus for **providing** data blocks describing three-dimensional terrain **to a render[er]**,³ the data blocks

² The italicized terms in the Claims quoted below were construed in the initial *Markman Memorandum*. *Skyline Software Sys. v. Keyhole, Inc.*, 421 F.Supp.2d 371 (D. Mass. 2006). The bold-face terms are construed in this Memorandum.

³ "Render" appears to be a typographical error and neither party contends that the term should be construed as anything

belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels, the apparatus comprising:

a local memory which stores data blocks corresponding to coordinates proximal to a current viewpoint of the renderer;

a communication link, through which the memory receives the data blocks from a remote server;

a processor which **receives** one or more specified coordinates along with indication of a respective resolution level **from a renderer, provides the renderer** with a *first data block* which includes data corresponding to the one or more specified coordinates from a local memory, and **downloads** over the communication link one or more data blocks of a resolution level higher than the resolution level of the first block which include data corresponding to the one or more coordinates **if the first block is not from the indicated level.**

(col. 18, ll. 12-31) (emphasis supplied).

From Claims 1 and 12, Defendants have asked for construction of the following additional terms: "downloading," "receiving [receives] from the renderer," "providing [provides] the renderer," and "downloading . . . if the provided block from the local memory is not at the indicated resolution level [if the first block is not from the indicated level]." Because each of these terms may have relevance in the context of the new claims, I have determined to construe each in turn, despite Skyline's objections.

1. Downloading

The term "downloading" is found throughout the '189 patent, which, at its core, concerns a method and apparatus for

other than "renderer."

transferring data structured in particular ways from one computer, generally a server, to another, generally a personal computer with limited storage capacity. Skyline fails to provide a specific proposed definition of "downloading" in either of its claim construction briefs, but Google characterizes Skyline's proposed construction as "transferring from a remote server to a local computer." Since this definition seems to comport with Skyline's understanding of the "plain and ordinary meaning" of the term discussed in its Reply brief and because the parties appear to have engaged in negotiations over disputed terms, I will accept this definition as an accurate reflection of Skyline's position. Google argues for a more specific and narrow definition, namely "requesting over a network and receiving in local memory from a separate computer."⁴

Google's proposed definition would add several limitations. First, it would make clear that nothing is "downloaded" until it has actually been received by the requesting computer. This limitation comports with the ordinary meaning of the term as it would have been understood by one skilled in the art, and it corresponds to the way the term is used at various points in the

⁴ At the November 1, 2006 Claim Construction Hearing, Google indicated a willingness to drop "local memory" in favor of simply requiring a download to the local computer. See Nov. 1, 2006 Transcript, p. 7:8-23; Defendants' Claim Construction Hearing Slides, p. 10 ("'Requesting over a network and receiving by the user's computer from a separate computer' is an acceptable and accurate construction.").

'189 patent.⁵ Use of the term "transferring" would, by contrast, be ambiguous as to whether the files transferred must necessarily be received by the target computer to be considered "downloaded." Because I find that the data must be received for the download to be complete, I adopt a construction of this element of the term that makes both "requesting" and "receiving" critical steps in the download process.

Google originally sought a limitation that the requested data blocks be received in "local memory" to be considered downloaded. As a matter of construction, this initially appeared to present a close question. Such a specific limitation is not necessarily imposed by the phrase "downloading," when used in a general context. One skilled in the art seems likely to have understood "downloading" as a phrase that required additional information, namely *where* the downloaded material would be sent (downloaded to a hard drive, etc.). The end point would in many cases be the local memory of a machine (RAM, cache memory, virtual memory, or other local memory options), but it could also be a disk drive, hard drive, or other storage device.

As used in the '189 patent, however, it is clear from the claims and specifications that "downloading" necessarily means

⁵ See, e.g., '189 patent at col. 12, ll. 5-7 ("cache manager preferably always requests that server send a block after the cache manager has received its parent block"); col. 12, ll. 66-67 ("[w]hen cache manager finishes downloading an additional block. . .") (emphasis supplied).

that the data blocks arrive in the "local memory." The clearest indication that this is the case is in Claim 12 (and repeated verbatim in Claims 13, 14, 15, 16, and 18), which describes the apparatus covered by the patent. The claimed apparatus has three parts: a local memory, a communication link, and a processor. The communication link is described as something "through which the memory receives the data blocks from a remote server" (col. 18, ll. 19-20) (emphasis supplied). The only reasonable construction of this phrase is that "the memory" refers to the "local memory" that is part of the apparatus described immediately above. The apparatus as described contains no other memory, so there is no non-local memory for the communication link to access. Consequently, the communication link must download into the local memory.

This reading is supported by Claim 7, discussed in more detail below, which describes "downloading . . . excess blocks . . . to fill up the local memory." (col. 17, ll. 58-60) No claim makes reference to downloading excess blocks for any other purpose, and Fig. 8 clearly indicates that excess blocks are discarded when the local cache memory is full, not moved to more permanent storage locations.⁶ Although "downloading" could

⁶ Skyline argues that Fig. 8 supports this position but I disagree. As noted in the patent, Fig. 8 "is a flow chart illustrating the actions of cache manager. . . ." The cache manager is described in the specification as something that preferably "manages a group of blocks and/or sub-blocks in a cache memory." (col. 11, ll. 39-40) Cache memory is "used [in the

theoretically mean something other than "downloading to the local memory," the fact remains that any time the term is qualified in the '189 patent, it specifies or strongly implies the use of local memory. Regardless, since both parties have agreed that "local memory" is not a necessary component of the claim construction, I will omit this term.

Construction: (downloading) Requesting over a network from a separate computer and receiving on a local computer

2. Receiving [receives] from the renderer; Providing

'189 patent] generally to refer to any relatively small memory which can be accessed rapidly by the processor and is used to save data which is most likely to be used by the processor." (col. 11, ll. 59-61) Based on this definition, cache memory is quintessentially local memory, because the entire purpose of local memory is to store often-used data for rapid access.

The description of Fig. 8 indicates that the cache manager starts by "downloading the first batch of level 1 blocks" and then moves into a wait state. Skyline reads this statement to imply that the initial download is received into something other than "local memory." However, this proposition fails to take into account the character of the cache manager. As Fig. 5 makes clear, the cache manager is part of the processor on the local machine, so any memory to which it has rapid access is fundamentally "local memory" as well, particularly as the term was initially construed. *Skyline Software Sys. v. Keyhole, Inc.*, 421 F.Supp.2d 371, 389 (D. Mass. 2006) ("Memory easily accessible to the user's processor . . . and distinct from the memory of the remote server from which data must be downloaded."). Skyline does not specify where it thinks the block in Fig. 8 is downloaded when it first enters the cache manager, but it seems clear that, wherever the downloaded block is initially held, it must be within the realm of local memory as previously construed. The apparatus descriptions in Claims 12, 13, 14, 15, 16, and 18 further reinforce the notion that the processor and cache manager are closely linked, by assigning tasks to the processor (providing data blocks to the renderer, etc.) that the cache manager is described as performing in various preferred embodiments.

[provides] the renderer

I construed the term "renderer" in the first *Markman* Memorandum as a "software and/or hardware object that performs at least the following functions: (1) determining and providing to another object the required coordinates in the terrain along with a respective resolution level; (2) receiving the data blocks corresponding to the specified coordinates; and (3) using the received data blocks to display a three-dimensional image."

Skyline Software Sys. v. Keyhole, Inc., 421 F.Supp.2d 371, 388 (D. Mass. 2006) (emphasis supplied). Google now argues that the phrases "receiving [receives] from the renderer" and "providing [provides] the renderer" should be construed to emphasize that something separate from the renderer either receives or provides the information handled by the renderer, depending on the context. Skyline simply maintains that these phrases should not be construed, and provides no alternate construction.

To the extent necessary to clarify the matter, I agree that the renderer must interact with other objects,⁷ whether hardware or software, and does not exist in a self-contained universe.

⁷ Skyline has expressed concern over the use of the word "object," which has a specific meaning in the context of computer programming. (Skyline's Claim Construction Slides, Nov. 1, 2006 at 13-15; Nov. 1, 2006 Transcript pp. 15:21-16:15). "Object," as used in this discussion, should not be read to refer specifically to software objects of the sort used in object-oriented programming. I use "object" in the more general sense of a "thing," which could be a software object, but need not be.

The '189 Patent is based on the conception of a renderer that receives information, over a communication link, and uses the transferred data to display three-dimensional terrain. In every description in the claims and specifications, and in the explanatory drawings, the renderer is depicted as receiving data from another location, and providing information to other discrete objects.⁸ This construction is consistent with the prior construction of "renderer," where I found a key function of the renderer to be "providing [information] to another object."⁹

Construction: (Receiving from the renderer; Receives from the renderer) Something distinct from the renderer receiving from the renderer; Something distinct from the renderer that receives from the renderer

Construction: (Providing the renderer; Provides the renderer)
Something distinct from the renderer providing to the renderer;
Something distinct from the renderer that provides to the renderer

3. **Downloading . . . if the provided block from the local memory is not at the indicated resolution level [if the first block is not from the indicated level]**

Google argues that the conditional nature of these phrases requires the requesting computer make a determination as to

⁸ See, e.g., col. 11, ll. 19-38 (describing the renderer).

⁹ *Skyline Software Sys. v. Keyhole, Inc.*, 421 F.Supp.2d at 388 (emphasis supplied).

whether the condition has been satisfied before taking any action, and proposes a construction of "downloading . . . upon a determination of whether the first data block is not of the indicated resolution level." Skyline proposes "downloading . . . when the first data block is not of the indicated resolution level."

Basic logic, as manifested in computer programming, makes clear that Google's argument must be correct. Although Skyline argues that the proposition "[w]hen an action is conditional, a computer program must make a determination as to whether the specified condition is met in order to take appropriate action" is not true, its rationale for this argument is unconvincing. Skyline attempts to read into "make a determination" a requirement that the computer test directly whether the conditional is satisfied. No such limitation is implied. The necessary determination may be made indirectly, by testing for a related condition, for example, or by forcing the conditional always to be satisfied, but these indirect tests do not eliminate the need to make *some* determination as to whether the conditional is satisfied before actions predicated on it are taken.¹⁰ Computers are logic-based machines -- they do not just "know"

¹⁰ Skyline acknowledged at the November 1, 2006 Claim Construction Hearing that some determination was necessary, but disagreed with the functional approach Google described. See Nov. 1, 2006 Transcript, pp. 19:1-20:7.

when a condition X is true without testing for X in some manner.

Therefore, I find that some determination, although not necessarily a direct test, is necessary before conditional actions are undertaken.

Construction: (Downloading . . . if the provided block from the local memory is not at the indicated resolution level)
Downloading . . . upon some determination that the block provided from local memory is not at the indicated resolution level

Construction: (Downloading . . . if the first block is not from the indicated level) Downloading . . . upon some determination that the first block is not at the indicated level

B. Disputed Terms in Claim 2

Claim 2, a dependent claim, reads:

A method according to claim 1, wherein downloading the one or more additional data blocks comprises downloading the blocks from a **succession of resolution levels**, from the level immediately higher than the resolution level of the first block up to the maximal existent resolution level on the server not above the indicated level.

(col. 16, ll. 45-50) (emphasis supplied).

1. Succession of resolution levels

The parties dispute one term in Claim 2, "succession of resolution levels." Google proposes a construction of "in order of increasing resolution level" to capture the hierarchical underpinnings of the claim, while Skyline argues that no construction is needed and does not propose an alternate construction.

When read in context, it is clear that the "succession" referred to is from lowest resolution to higher resolution. The phrase in question is directly followed by instructions to move "*from the level immediately higher than the resolution level of the first block up to the maximal existent resolution level on the server not above the indicated resolution level.*" The process seems quite clear: an initial block (generally of fairly low resolution) is acquired, a new block is downloaded from the next highest resolution level, then additional higher resolution blocks are downloaded until an upper limit (either the highest resolution level stored in the database, or the requested top resolution level) is reached. Claim 2 specifically indicates that this process happens from low resolution to high resolution, in succession. The patent history also supports this reading, as Skyline distinguished the '189 Patent from prior art based on the ordered nature of the download. (See February 26, 1999 Amendment, Goog 000118, GOOG 000153) Consequently, I find that the blocks referenced in Claim 2 are, by definition, downloaded "in order of increasing resolution level."

Construction: (succession of resolution levels) In order of increasing resolution level

C. Disputed Terms in Claim 3¹¹

Claim 3 reads:

A method of providing data blocks describing three-dimensional terrain to a renderer, the data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels, the method comprising:

receiving from the renderer a plurality of coordinates in the terrain along with indication of a respective resolution level; said **plurality of coordinates being included in a plurality of respective distinct blocks**;

providing the renderer with first data block which includes data corresponding to at least some of the plurality of coordinates from a local memory;

downloading from a remote server one or more additional blocks which include data corresponding to a plurality of respective distinct blocks if the provided block from the local memory is not at the indicated resolution level, wherein blocks of lower resolution levels are downloaded before blocks of higher resolution levels.

(col. 16, ll. 51-67) (emphasis supplied).

1. Plurality of coordinates being included in a plurality of respective distinct blocks

The focus of the parties' dispute over the construction of Claim 3 concerns whether there is necessarily a one-to-one correspondence between one set of coordinates (out of the several sets provided) and a particular block in the set of data blocks. Google proposes a construction of "each one of the plural sets of coordinates being included in a separate distinct one of a plurality of data blocks describing three-dimensional terrain," while Skyline proposes "more than one set of coordinates being described by the data contained in more than one data block."

¹¹ Identical terminology is found in Claims 4 and 5.

Comparing Claim 3 to Claim 1, it is clear that Claim 1 is not limited to situations where a single point is provided to the renderer. Rather, the renderer may receive "one or more coordinates," indicating that several coordinates, defining particular points, could be provided to the renderer simultaneously. The question then necessarily arises how Claim 3, which calls for several sets of coordinates to be passed to the renderer, differs from Claim 1. The difference is illuminated in the detailed description of Fig. 7, which describes two different ways the renderer might specify which blocks are necessary.

Preferably, renderer determines the exact blocks needed and calls for them using their (x,y) coordinates and their resolution level. Alternatively or additionally, renderer specifies, for each resolution level, the coordinates of the boundaries of the necessary areas, and cache manager determines the identities of the required blocks.

(col. 14, ll. 10-16).

The method described in Claim 1 seems to correspond to the first method in the Fig. 7 description (providing specific required points), while the Claim 3 method appears to follow the second description (using several points to define the boundaries of an area). Although Fig. 7 is not itself part of a claim, of course, it does shed light on the meaning of Claim 3, particularly given the somewhat Delphic phrase at issue.

As Skyline notes (Plaintiff's Claim Construction Slides, p. 22), "respective distinct" modifies blocks, not the plurality of

coordinates. However, Skyline's proposed construction reads "respective distinct" out of the claim entirely. I read the disputed language in Claim 3 as largely descriptive -- the several sets of coordinates at issue are, by definition, contained within available data blocks, which are themselves distinct from one another. In this reading, several sets of coordinates may be contained within one distinct block, but the blocks do not overlap. In this way, the blocks are "distinct," and coordinates are found in their "respective" blocks. Because I do not find anything in the claim language or elsewhere requiring a strict one-to-one correspondence between coordinates and blocks (so that each point describes only one data block, and the same number of points and blocks are contained in each plurality), I construe this language more broadly, allowing a one-to-many relationship between the blocks and the coordinates,¹² while still emphasizing the distinctiveness of each block.

Construction: (plurality of coordinates being included in a plurality of respective distinct blocks) Several coordinates, where each set of coordinates is contained within one block in a set composed of data blocks that are distinct from one another

¹² The concept is that one block may contain many coordinates and each set of coordinates may be contained within only one block. This one-to-many concept is in contrast to Google's one-to-one and Skyline's many-to-many.

D. Disputed Terms in Claim 7¹³

Claim 7 reads:

A method of providing data blocks describing three-dimensional terrain to a renderer, the data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels, the method comprising:
receiving from the renderer one or more coordinates in the terrain along with indication of a respective resolution level;
providing the renderer with a first data block which includes data corresponding to the one or more coordinates, from a local memory;
downloading from a remote server one or more additional data blocks which include data corresponding to the one or more coordinates if the provided block from the local memory is not at the indicated resolution level; and
downloading from a remote server excess blocks not currently needed by the renderer to fill up the local memory **when not downloading blocks required by the renderer.**

(col. 17, ll. 43-63) (emphasis supplied).

1. When not downloading blocks required by the renderer

The parties dispute two elements of this phrase. First, whether it is meaningful to construe "when" as indicative of a period of time, and, second, what it means for a block to be "required" by the renderer (whether this means only blocks necessary for the current display, or whether it includes other blocks requested by the renderer that are not required for the immediate view). Google proposes a definition of "during periods of time when the local computer is not downloading data blocks

¹³ Identical terminology is found in Claim 18.

describing three-dimensional terrain in response to the coordinates received from the renderer." Skyline counters with "when not downloading data for displaying the scene corresponding to the current view."

Skyline argues that Google's use of "during periods of time when the local computer is not downloading" is problematic, because computers tend to run processes on multiple threads simultaneously, and download data on multiple simultaneous connections, meaning the entire computer is rarely at rest (even if one thread or connection might be free to download additional data). There is some merit to this suggestion, but there are clearly periods of time in which individual threads and connections are not downloading requested data, and are free to download excess blocks for the local memory. Therefore, I adopt Google's construction of a "period of time," but emphasize that the entire computer need not be at rest before excess data blocks are downloaded.

Skyline next argues that the only blocks "required by the renderer" are those immediately necessary for the display of the current view. This proposed construction overstates the point. Clearly there are at least two types of data blocks -- those required by the renderer, and those that are being downloaded to fill up excess local memory, with the expectation that they might eventually be required for display of a particular viewpoint. To say that those "required by the renderer" are only those

immediately needed for the current viewpoint is not wholly accurate.

Because the data blocks are not transferred at the exact moment they are requested (due to network lag, processor delays, etc.), there will often be a backlog of data blocks that the renderer has requested, but not yet received. Cutting down on this time lag is a purpose of the '189 patent. However, because the lag exists, and the '189 patent allows for a constantly changing user viewpoint, there will often be a queue of data blocks, some of which may not be needed any longer, that have been requested by the renderer, but not yet downloaded. Until these requested blocks have finished downloading, excess blocks will not be downloaded to fill the local memory. See Fig. 8 ("Queue Empty?" No --> Download blocks from queue;¹⁴ Yes --> Download excess surrounding blocks). Therefore, files which are, or were, "required by the renderer" may include data blocks not required for the currently displayed view.¹⁵

¹⁴ The process diagrammed in Fig. 8 does run a check to see if the block in the queue is out of range (which it would be if the viewpoint has changed so dramatically that the block is obviously no longer needed), but excess blocks are not automatically downloaded if the requested block is found to be out of range. Instead, the logic says to go back to the top of the flowchart, and wait for further inputs.

¹⁵ This point becomes particularly clear if one envisions a viewer hovering over a point, say at the top of a hill. The viewer might start by looking north, then turn around in a circle, examining each of the cardinal points. By the time the viewer's attention comes back to the starting point, the resolution of the terrain image should have increased, so that

Construction: (when not downloading blocks required by the renderer) During periods of time when the local computer, or a connection thereof, is not downloading data blocks¹⁶ in response to coordinates received from the renderer

E. Disputed Terms in Claim 8 and Claim 22

Claim 8 reads: "A method according to claim 7, wherein downloading the data blocks comprised downloading the blocks via the **Internet**." (col. 17, ll. 62-64) Similarly, Claim 22 reads: "Apparatus according to claim 18, wherein the communication link comprises a connection to the **internet**." (col. 20, ll. 37-38) (emphasis supplied)

1. Internet

The parties appear to attach no significance to the fact that the disputed term is capitalized in one instance and not in another, so I will assume "Internet" is a proper noun in both cases. The crux of the dispute between the parties over this term is whether private networks are specifically excluded from "the Internet." Skyline argues that the Internet excludes

the view is more detailed than it was initially. While the viewer was turning around, data blocks "required by the renderer" to improve resolution at the starting point were being downloaded. However, these blocks were not required for the display when the viewer was looking in other directions (due south, for example).

¹⁶ I have eliminated Google's proposed addition of "describing three-dimensional terrain" whenever "data blocks" appeared, because this term has already been construed in a way that makes the additional language unnecessary.

"private networks even if they use internet protocols or have connections to the Internet." I find this exclusion too broad, but agree that *some* private networks may not be part of the Internet.

In ordinary usage, the Internet refers to the communication network composed of a multitude of interconnected smaller networks, which communicate via Internet Protocols. Some of the smaller networks are fully public, some are private, and most fall somewhere in between. An argument that "private networks" are not part of the Internet would be surprising to any home user who has a password protected wireless network. Clearly the home wireless network is *private*, in that only authorized users may log on, but it is also a part of the Internet. In fact, the Internet is composed largely of interconnected private networks, where access to any particular network is controlled via passwords or other authentication techniques.

However, there does seem to be a distinction between a private, in the sense of authenticated, network, and a wholly self-contained group of computers or other devices communicating only with one another. All parties in the November 1, 2006 Claim Construction hearing agreed that three computers hardwired to each other, and communicating only among themselves, were probably not part of "the Internet," even if they communicated via Internet Protocol. See Nov. 1, 2006 Transcript at p. 75:19-

22. Although the public nature of the Internet should be captured in the idea of a "publicly accessible network," I will include a specific limitation on wholly self-contained private networks, to clarify that these are not part of "the Internet." However, if the private network is not wholly self-contained, in that it has a link to the broader Internet, I find that the network is part of the Internet, at least for the purposes of Claim 22 (which only requires a "connection to the internet"). For Claim 8, at least some portion of the download must happen over the broader Internet for the claim language to be satisfied. In general, I also agree with Skyline that the Internet is a single entity (the Internet), not several (an Internet).¹⁷

Construction: (Internet) The publicly accessible network capable of relaying information via Internet Protocol, either alone or in conjunction with one or more other protocols, but not including a wholly self-contained private network of devices communicating only with each other

F. Disputed Terms in Claim 9

Claim 9 reads:

A method according to claim 7, wherein the

¹⁷ However, this construction is not intended to exclude networks such as Intenet2, which are built on separate physical infrastructure, but are essentially updated and experimental versions of the current Internet. Should more than one mesh of networks come to have the characteristics of "the Internet" as described, the term should be read to describe all such networks.

renderer renders a view from a current viewpoint, and wherein downloading the excess blocks comprises filling the local memory with **substantially all of the blocks surrounding a point in the terrain seen from the current viewpoint within a predetermined distance range.**

(col. 17, l. 65 - col. 18, l. 3) (emphasis supplied).

1. Substantially all of the blocks surrounding a point in the terrain seen from the current viewpoint within a predetermined distance range

The parties dispute two elements of this phrase: first, the meaning of "surrounding" and whether this word implies a three-dimensional sphere and a uniform distance from a point, and, second, whether the original viewpoint is included in the set of "points" referred to in Claim 9. Google proposes a construction of "substantially all of the excess data blocks describing three-dimensional terrain on all sides (in all directions) out to a pre-established distance from a point in the terrain that is seen from the current viewpoint," while Skyline suggests "substantially all of the blocks which include data covering terrain which is within a predetermined distance range in one or more directions from either the viewpoint or a point in the terrain visible from the current viewpoint."

a. Surrounding

To define "surrounding," Google suggests that the plain meaning of the word is to "enclose on all sides." The ambiguity latent in the term in this context is that the rendered display

is somewhere between a two-dimensional and a three-dimensional object, so it is unclear whether a circle or a sphere would "surround" a point in the terrain. As Skyline correctly points out, Google's construction would imply that underground blocks of data would have to be downloaded for the "surrounding" data blocks to be considered gathered. It seems unlikely that this result was intended by the claim, because the '189 patent is primarily concerned with the rendering of surfaces with a minimal vertical dimension, not true three-dimensional terrain. Consequently, I find that "surrounding" does not imply a sphere surrounding a point on all sides, but a circle surrounding a point in two dimensions.

The next disputed issue with respect to "surrounding" is whether the area of excess blocks to be downloaded must necessarily be uniform in all directions (or whether a wider area of blocks could be pre-selected in the direction of travel, potentially creating an ellipse around a reference point, rather than a circle).¹⁸ "Surrounding" is customarily read as encircling in a uniform manner,¹⁹ and numerous references in the

¹⁸ I also note that the points in the terrain selected as the starting points in Claim 9 could be preferentially chosen based on the direction of travel. Although the blocks surrounding each point would be downloaded to a uniform distance, the points chosen could be weighted towards the direction of travel.

¹⁹ See, e.g., Webster's Third New Int'l Dictionary (Unabridged) (1986) (listing as synonyms "encircle," "circle," "ring," and "encompass").

'189 patent support this reading.²⁰ At several points, the '189 patent mentions that the user can control the direction of view and look around in all directions, which necessitates downloading blocks in all directions.²¹ Based on the common understanding of "surrounding," the specific use of the term and concept in the '189 patent, and the broader context of the claimed method and apparatus, I construe "surrounding" to mean all of the blocks within a uniform distance from the point in question.

b. Point in the terrain

The next dispute in the construction of this phrase is whether the "current viewpoint" is a "point in the terrain seen from the current viewpoint," as Skyline argues. The plain language of the phrase makes clear that it is not. The current

²⁰ See, e.g., col. 12, ll. 23-25 ("blocks, which are most preferably organized in a square centered directly below the location of the viewpoint"); col. 16, ll. 4-6 ("cache manager first downloads the eight blocks surrounding the block which is directly below the current viewpoint").

²¹ See, e.g., col. 2, ll. 3-5 ("allowing the pilot to see the view seen at any point along the flight course at substantially any desired angle"); col. 2, ll. 19-21 ("A user may select at substantially each point along the route the direction of view and may change the direction dynamically"); col. 7, ll. 6-7 ("Preferably, the viewpoint is controlled by a user of the processor"); col. 11, ll. 9-10 ("there is no compulsory correlation between the flight direction and the view direction"); col. 15, ll. 63-67 ("if the queue is empty, cache manager fills cache memory with the blocks within the range of the current viewpoint, so that, for any direction of view from the current viewpoint, there is no need to download further blocks from the server") (emphasis supplied).

viewpoint is the current viewpoint. A point seen *from* the current viewpoint is inherently different from the viewpoint itself. Consequently, I reject Skyline's argument that the current viewpoint is also a point seen from the current viewpoint.

Finally, Google includes "excess blocks" in its construction. This term is unnecessary, as Claim 9, a dependent claim of Claim 7, only covers the download of excess blocks.

Construction: (substantially all of the blocks surrounding a point in the terrain seen from the current viewpoint within a predetermined distance range) Substantially all of the data blocks covering terrain within a uniform predetermined distance from a point in the terrain that is seen from the current viewpoint

/s/ Douglas P. Woodlock

DOUGLAS P. WOODLOCK
UNITED STATES DISTRICT JUDGE

Appendix

Disputed Term	Court Construction
downloading	requesting over a network from a separate computer and receiving on a local computer
receiving (receives) from the renderer	something distinct from the renderer receiving from the renderer something distinct from the renderer that receives from the renderer
providing (provide) the renderer	something distinct from the renderer providing to the renderer something distinct from the renderer that provides to the renderer
downloading . . . if the provided block from the local memory is not at the indicated resolution level <u>or</u> downloads . . . if the first block is not from the indicated level	downloading . . . upon some determination that the block provided from local memory is not at the indicated resolution level downloads . . . upon some determination that the first block is not at the indicated resolution level
succession of resolution levels	in order of increasing resolution level
plurality of coordinates being included in a plurality of respective distinct blocks	several coordinates, where each set of coordinates is contained within one block in a set composed of data blocks that are distinct from one another
when not downloading blocks required by the renderer	during periods of time when the local computer, or a connection thereof, is not downloading data blocks in response to coordinates received from the renderer
Internet	the publicly accessible network capable of relaying information via Internet Protocol, either alone or in connection with one or more other protocols, but not including a wholly self-contained private network of devices communicating only with each other
substantially all of the blocks surrounding a point in the terrain seen from the current viewpoint within a predetermined distance range	substantially all of the data blocks covering terrain within a uniform predetermined distance from a point in the terrain that is seen from the current viewpoint